

A regularizing iterative ensemble Kalman method for PDE-constrained inverse problems

Marco A. Iglesias

School of Mathematical Sciences, The University of Nottingham, University Park,
Nottingham, NG7 2RD, United Kingdom.

Abstract. We introduce a derivative-free computational framework for approximating solutions to nonlinear PDE-constrained inverse problems. The general aim is to merge ideas from iterative regularization with ensemble Kalman methods from Bayesian inference to develop a derivative-free stable method easy to implement in applications where the PDE (forward) model is only accessible as a black box (e.g. with commercial software). The proposed regularizing ensemble Kalman method can be derived as an approximation of the regularizing Levenberg-Marquardt (LM) scheme [14] in which the derivative of the forward operator and its adjoint are replaced with empirical covariances from an ensemble of elements from the admissible space of solutions. The resulting ensemble method consists of an update formula that is applied to each ensemble member and that has a regularization parameter selected in a similar fashion to the one in the LM scheme. Moreover, an early termination of the scheme is proposed according to a discrepancy principle-type of criterion. The proposed method can be also viewed as a regularizing version of standard Kalman approaches which are often unstable unless ad-hoc fixes, such as covariance localization, are implemented.

The aim of this paper is to provide a detailed numerical investigation of the regularizing and convergence properties of the proposed regularizing ensemble Kalman scheme; the proof of these properties is an open problem. By means of numerical experiments, we investigate the conditions under which the proposed method inherits the regularizing properties of the LM scheme of [14] and is thus stable and suitable for its application in problems where the computation of the Fréchet derivative is not computationally feasible. More concretely, we study the effect of ensemble size, number of measurements, selection of initial ensemble and tunable parameters on the performance of the method. The numerical investigation is carried out with synthetic experiments on two model inverse problems: (i) identification of conductivity on a Darcy flow model and (ii) electrical impedance tomography with the complete electrode model. We further demonstrate the potential application of the method in solving shape identification problems that arises from the aforementioned forward models by means of a level-set approach for the parameterization of unknown geometries.

Submitted to: *Inverse Problems*

1. Introduction

We propose a computational derivative-free regularization method for approximating solutions to nonlinear PDE-constrained inverse problems. More precisely, the aim of the

method is to identify parameters in PDE models given data/observations of the solution of the PDE. Inverse problems of this kind are ill-posed in the sense of stability; small perturbations of the data may have uncontrolled effects on the approximation of the unknown parameters in the PDE. Therefore, the computation of stable approximations of solutions to these inverse problems requires *regularization*. Classical regularization methods (e.g. Tikhonov regularization) [12] reformulate the inverse problem so that the regularized version can be solved with, for example, standard optimization methods. In contrast, *iterative regularization* approaches regularize while computing a stable approximation to the inverse problem [22]. The aim of these methods is to compute an estimate/approximation controlled by the noise level. As the noise in the data goes to zero, this approximation converges to a solution of the identification problem.

1.1. Contribution of this work

Most existing iterative regularization methods [22] require the implementation of the Fréchet derivative of the forward map as well as the corresponding adjoint operator. In various applications, however, standard software for forward simulation does not provide numerical approximations of the Fréchet derivative and/or its associated adjoint. In some cases, even if the linearization of the forward map is computed within the forward simulation (e.g. in a Newton-type solver for nonlinear PDEs) this may not be accessible in a modular fashion suitable for an iterative regularization framework. In this paper we present a derivative-free ensemble Kalman-based iterative regularization technique for the approximation of PDE-constrained identification problems. While standard ensemble Kalman methods are aimed at approximating an inverse problem posed in a Bayesian inference framework, the objective of the present work is to merge ideas from iterative regularization with ensemble Kalman methods to develop a computational framework with the regularization properties needed to solve classical (deterministic) identification problems. The proposed framework offers the flexibility of typical implementations of ensemble Kalman methods often used for large-scale data assimilation. In particular, it uses the forward map in a black-box fashion which makes it easy to implement and thus ideal for applications where modeling and simulation are performed with complex computer codes.

The proposed method can be derived as an approximation of the regularizing Levenberg-Marquard (LM) scheme developed by Hanke in [14]. More specifically, we construct an iterative ensemble method from the regularizing LM scheme by replacing operators involving the Fréchet derivative of the forward map and its adjoint by empirical covariances from an ensemble of elements from the parameter space. Members from this ensemble are iteratively updated according to an expression that resembles the Kalman update from standard ensemble Kalman filter/smoothing methods [13]. However, in contrast to those standard methods, we propose an update formula with a regularization parameter whose selection is made similar to the one of the regularization parameter in the LM scheme [14] but with the aforementioned ensemble approximations of the

forward map and its adjoint. Similarly, we propose an early termination of the scheme motivated by the discrepancy principle used in iterative regularization methods. The proposed ensemble Kalman regularizing scheme can then be regarded as (i) a derivative-free approximation of the regularizing LM scheme and/or (ii) a regularizing version of standard iterative Kalman methods [16].

An iterative regularizing ensemble Kalman method of the type presented here has been recently introduced in [19] for the solution of Bayesian inverse problems in reservoir modeling applications. The aim of [19] was to show that importing ideas from iterative regularization may improve the performance of Kalman methods for approximating the Bayesian posterior. The classical approach to the inverse problem is now pursued in the present work. More concretely, we wish to assess the convergence and regularizing properties of ensemble Kalman methods derived from iterative regularization techniques. In contrast to [19] where the framework was Bayesian and focused on reservoir applications, our objective here is to study the performance of regularizing ensemble methods for solving classical identification problems in generic PDE-constrained applications.

The contribution of the present work is threefold. First, we show that the proposed method addresses the ill-conditioning typically exhibited by existing implementations of ensemble methods. We show that importing ideas from iterative regularization such as the discrepancy principle can offer stability that, in existing implementations are often treated with ad-hoc methodologies such as covariance localization and covariance inflation. The second contribution of this paper is to showcase the potential application of the proposed methods for addressing a wide range of parameter identification problems. We show that, with reasonable computational cost, the proposed method can be not only computationally advantageous but also robust and accurate. We consider two model inverse problems: (i) identification of hydraulic head in a Darcy flow model and (ii) electrical impedance tomography (EIT) with a complete electrode model (CEM). In addition, we display the capabilities of the proposed method to solve shape identification problems where the computation of the shape derivative of the forward map may be cumbersome. In concrete, we combine the proposed ensemble method with the level-set approach of [17] to estimate shapes whose boundaries determine regions of sharp discontinuities of parameters in the PDE models under consideration. Since our methodology does not require derivatives, our level-set based formulation is considerably more simple than standard approaches where shape derivatives are needed [6, 5, 30]. Moreover, the proposed ensemble level-set approach does not use the level-set equation as in standard formulations; our iterative scheme induces a stable evolution of the unknown interface. In addition, by initializing the ensemble according to the ideas recently proposed in [17] we avoid computational issues such as the flattening of the level-set function which is often observed with standard methods and typically addressed with ad-hoc techniques. The third contribution of the present work is to provide an extensive numerical investigation in order to assess the convergence and regularization properties of the proposed ensemble scheme with respect to ensemble

size, number of measurements, initial ensemble and tunable parameters. Although the convergence theory of the proposed regularizing ensemble Kalman method is an open problem, our numerical study offers deep insight into the potential application of iterative regularization for the development of derivative-free ensemble methods thereby opening up a whole new field for application and theory.

The rest of the manuscript is organized as follows. In Section 2 we introduce the general framework for the identification problem. The proposed ensemble Kalman method is introduced in subsection 2.2. The test models used for the validation of the proposed scheme are introduced in subsection 2.3. Preliminary numerical examples that show the regularizing properties of the method are presented in subsection 2.4. In Section 3 we discuss general aspects and properties of the proposed scheme including the derivation of the method as an ensemble approximation of the regularizing LM scheme of [14]. In Section 4 we then conduct an extensive numerical investigation of the proposed scheme. In particular, we study the relation between the ensemble size and the number of measurements in terms of their effect on the regularization and convergence properties of the scheme. We investigate the effect of the tunable parameters of the scheme and the regularization properties with respect to the ensemble size in the small noise limit. In Section 5 we discuss some potential applications of the scheme for the solution of geometrical inverse problems. Conclusions and future research are provided in Section 6.

2. PDE-constrained inverse problems

2.1. General framework

Let us denote by $G : X \rightarrow Y$ the (nonlinear) *forward operator* that arises from the PDE-constrained model under consideration. In other words, G maps the space X of PDE parameters (e.g. coefficients, source terms and/or boundary conditions) to the observation space Y defined in terms of observable quantities related to the solution of the PDE (e.g. pointwise measurements of the solution). We assume that X and Y separable Hilbert spaces with norms denoted by $\|\cdot\|_X$ and $\|\cdot\|_Y$ respectively. For the applications under consideration, the dimension of the observation space is often small (i.e. order of $10^3 - 10^4$). Therefore, we consider the case where Y is finite dimensional. However, upon discretization, the dimensions of X could be large (i.e. greater than 10^6). Therefore, in order to derive algorithms robust under grid refinement we consider the case where $\dim(X) = \infty$; our aim is thus to keep our formulation within a functional analytical framework.

The PDE-constrained model has an unknown parameter $u^\dagger \in X$ that we wish to identify from noisy measurements $y^\eta \in Y$ defined by

$$y^\eta \equiv G(u^\dagger) + \xi \tag{1}$$

where $\xi \in Y$ is noise which for the purpose of this exposition will be considered deterministic. In addition to y^η , we assume we have knowledge of the (weighted) noise

level η defined by

$$\eta \equiv \|\Gamma^{-1/2}(y^\eta - G(u^\dagger))\|_Y \quad (2)$$

where $\Gamma : Y \rightarrow Y$ is a self-adjoint positive-definite operator which for the present analysis can be understood as a weighting/scaling operator that enable us, for example, to include information concerning the precision of our measurement device.

The identification problem is the following: *Given y^η , find $u \in X$ such that $G(u) = y^\eta$.* Since the data may not be in the range of the forward operator we may alternatively formulate the identification as the minimizer of

$$\Phi(u) \equiv \|\Gamma^{-1/2}(y^\eta - G(u))\|_Y \quad (3)$$

However, forward operators G that arise from PDE models are typically compact and weakly sequentially closed [11]. From this property it follows that the inverse problem is ill-posed in the sense of stability. In other words, we may find u_n such that $G(u_n) \rightarrow y^\eta$ but $u_n \not\rightarrow u^\dagger$. Therefore the computation of a minimizer of (3) with standard (unregularized) iterative minimization approaches may be unstable. This lack of stability can be alleviated by means of iterative regularization methods. As discussed in Section 1, iterative regularization provides a computational framework to compute stable approximation to the inverse problem. In concrete, their aim is to compute an estimate/approximation u^η controlled by the noise level η , that converges, in the small noise limit, to a solution of the identification problem, i.e. $u^\eta \rightarrow u$ as $\eta \rightarrow 0$, where the limit u satisfies $G(u) = G(u^\dagger)$. Most iterative regularization approaches require the computation of the Fréchet derivative of the forward operator G . In some applications where commercial software is used for forward simulations Fréchet derivatives may not be available and so the application of those iterative methods may be limited. The objective of the present work is to introduce a derivative-free regularization ensemble Kalman method for the stable computations of the identification problem.

2.2. The regularizing ensemble Kalman method

Let us assume that we are given an ensemble of N_e elements $u_0^{(j)}$ ($j \in \{1, \dots, N_e\}$) of the parameter space X . One may think of $\{u_0^{(j)}\}_{j=1}^{N_e}$ as an ensemble of potential initial guesses for any iterative regularization method applied for the stable identification of the unknown parameter u^\dagger . We will discuss the selection of such initial ensemble $\{u_0^{(j)}\}_{j=1}^{N_e}$ in subsection 3.1.

We now propose an iterative scheme where, at every iteration level n , each ensemble member $u_n^{(j)}$ is updated in such a way that the corresponding ensemble mean

$$\bar{u}_n \equiv \frac{1}{N_e} \sum_{j=1}^{N_e} u_n^{(j)} \quad (4)$$

approximates the inverse problem in the small noise limit as described in subsection 2.1. In other words, we need that algorithm stops at finite iteration level n^* producing an estimate $u^\eta \equiv \bar{u}_n^*$ such that $u^\eta \rightarrow u$ as $\eta \rightarrow 0$, where $G(u) = G(u^\dagger)$. The proposed regularizing ensemble Kalman method is presented below.

Algorithm 1 *Iterative regularizing ensemble Kalman method*

Let $\{u_0^{(j)}\}_{j=1}^{N_e} \subset X$ be the initial ensemble of N_e elements. Let $\rho \in (0, 1)$ and $\tau > 1/\rho$. For $n = 0, 1, \dots$

(1) **Prediction step.** Evaluate

$$w_n^{(j)} = G(u_n^{(j)}), \quad j \in \{1, \dots, N_e\} \quad (5)$$

and define $\bar{w}_n = \frac{1}{N_e} \sum_{j=1}^{N_e} w_n^{(j)}$

(2) **Discrepancy principle.** If

$$\|\Gamma^{-1/2}(y^\eta - \bar{w}_n)\|_Y \leq \tau\eta \quad (6)$$

stop. Output $\bar{u}_n \equiv \frac{1}{N_e} \sum_{j=1}^{N_e} u_n^{(j)}$.

(3) **Analysis step.** Define C_n^{uw} , C_n^{ww} by

$$C_n^{ww}(\cdot) = \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (G(u_n^{(j)}) - \bar{w}_n) \langle G(u_n^{(j)}) - \bar{w}_n, \cdot \rangle_Y \quad (7)$$

$$C_n^{uw}(\cdot) = \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (u_n^{(j)} - \bar{u}_n) \langle G(u_n^{(j)}) - \bar{w}_n, \cdot \rangle_Y. \quad (8)$$

Update each ensemble member:

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw}(C_n^{ww} + \alpha_n \Gamma)^{-1}(y^\eta - w_n^{(j)}), \quad j \in \{1, \dots, N_e\} \quad (9)$$

where α_n is chosen by the following sequence

$$\alpha_n^{i+1} = 2^i \alpha_n^0. \quad (10)$$

where α_n^0 is an initial guess. We then define $\alpha_n \equiv \alpha_n^N$ where N is the first integer such that

$$\alpha_n^N \|\Gamma^{1/2}(C_n^{ww} + \alpha_n^N \Gamma)^{-1}(y^\eta - \bar{w}_n)\| \geq \rho \|\Gamma^{-1/2}(y^\eta - \bar{w}_n)\|$$

Note that, in contrast to other regularization techniques, the discrepancy principle in (6) is applied to \bar{w}_n which (see expression (26)) is the average of the model output $G(u^{(j)})$ of each ensemble member. While this quantity approximates $G(\bar{u}_n)$ to first order (see subsection 3.3), it is clearly not the data misfit of the proposed estimate.

The proof of convergence of Algorithm 1 to a stable solution of the inverse problem is beyond the scope of the present manuscript; our aim is to offer numerical evidence of the convergence and regularization properties of the scheme.

2.3. Test Models

The proposed ensemble Kalman method will be tested on two PDE-constrained inverse problems. We consider small test inverse problems where we have the computational flexibility to conduct a large amount of numerical experiments in order to understand the effect that the tunable parameters, ensemble size, selection of initial ensemble and number of measurements have on the regularization properties and accuracy of the proposed method. We introduce the test models under consideration below.

2.3.1. Test model I. Darcy flow. We consider single-phase steady-state Darcy flow in a two-dimensional confined aquifer whose physical domain is $D = [0, 6] \times [0, 6]$. κ denotes the hydraulic conductivity. The flow is described in term of the piezometric head $h(x)$ ($x \in D$) given by the solution to [3]

$$-\nabla \cdot \kappa \nabla h = f \quad \text{in } D \quad (11)$$

where f is the source which for the present work is defined by

$$f(x_1, x_2) = \begin{cases} 0 & \text{if } 0 < x_2 \leq 4, \\ 137 & \text{if } 4 < x_2 < 5, \\ 274 & \text{if } 5 \leq x_2 < 6. \end{cases} \quad (12)$$

The following boundary conditions are considered

$$h(x, 0) = 100, \quad \frac{\partial h}{\partial x}(6, y) = 0, \quad -\kappa \frac{\partial h}{\partial x}(0, y) = 500, \quad \frac{\partial h}{\partial y}(x, 6) = 0, \quad (13)$$

We are interested in recovering the logarithm of the hydraulic conductivity $u \equiv \log \kappa$, from noisy pointwise measurements of the piezometric head h . In other words, we consider

$$G(u) = (h(x_1), \dots, h(x_M)) \quad (14)$$

where h is the solution to (11)-(13) and $\{x_j\}_{j=1}^M$ are the measurement locations. This groundwater model was first used as benchmark for inverse modeling in [7]. It has been also used as a test model for the identification of parameters with iterative regularization methods in [14, 15] and with an ensemble Kalman approach in [16].

2.3.2. Test model II. Complete Electrode Model. Our second test model is based on the Complete Electrode Model (CEM) for Electrical Impedance Tomography (EIT). The objective of EIT is to identify the conductivity κ of a body D given measurements of voltages from a configuration of m_e electrodes on $\{e_k\}_{k=1}^{m_e}$ places on the boundary ∂D . The measured voltage arise from current patterns applied on those electrodes. The forward model associated to EIT is the CEM which consist of computing the voltage v in D and the voltages $\{V_k\}_{k=1}^{m_e}$ on $\{e_k\}_{k=1}^{m_e}$ that satisfy

$$\nabla \cdot \kappa \nabla v = 0 \quad \text{in } D, \quad (15)$$

$$v + z_k \kappa \nabla v \cdot \mathbf{n} = V_k \quad \text{on } e_k, \quad k = 1, \dots, m_e, \quad (16)$$

$$\nabla v \cdot \mathbf{n} = 0 \quad \text{on } \partial D \setminus \cup_{k=1}^{m_e} e_k, \quad (17)$$

$$\int_{e_k} \kappa \nabla v \cdot \mathbf{n} \, ds = I_k \quad k = 1, \dots, m_e, \quad (18)$$

where $\{I_k\}_{k=1}^{m_e}$ are the currents injected through the electrodes, $\{z_k\}_{k=1}^{m_e}$ are the contact impedances of the electrodes and $u = \log \kappa$ where κ is the conductivity. Well posedness of the CEM model requires conservation of charge

$$\sum_{k=1}^{m_e} I_k = 0$$

Given κ and $\{z_k\}_{k=1}^{m_e}$, for each current patter $I = \{I_k\}_{k=1}^{m_e}$ there exists a unique solution $[v, \{V_k\}_{k=1}^{m_e}]$ to the CEM [31]. The EIT problem consists now of finding κ and $\{z_k\}_{k=1}^{m_e}$ from a set of n_p measurements of voltages $V_1 = \{V_{1,k}\}_{k=1}^{m_e}, \dots, V_{n_p} = \{V_{n_p,k}\}_{k=1}^{m_e}$ obtained from n_p current patters $I_1 = \{I_{1,k}\}_{k=1}^{m_e}, \dots, I_{n_p} = \{I_{n_p,k}\}_{k=1}^{m_e}$. For simplicity we assume that the contact impedances of the electrodes are known. Therefore, the identification problem is to find κ given

$$G(u) = [\{V_{1,k}\}_{k=1}^{m_e}, \dots, \{V_{n_p,k}\}_{k=1}^{m_e}]$$

For a review of the EIT problem we refer the reader to [4].

2.4. Regularizing properties of the ensemble Kalman method.

Although ensemble Kalman methods have been typically used in the statistical Bayesian framework, several publications [16, 24] have explored the use of these approaches for solving (deterministic) inverse problems such as the parameter identification PDE-constrained problems described earlier. However, most of these approaches are based on update formulas of the form of (9) but with a fixed parameter $\alpha_n = 1$, i.e.

$$u_{n+1}^{(j)} = u_n^{(j)} + C_n^{uw} (C_n^{ww} + \Gamma)^{-1} (y - w_n^{(j)}) \quad (19)$$

As we discuss in subsection 3.5, the standard choice $\alpha_n = 1$ is motivated by the application of Kalman methods for solving Bayesian inference problems when the model G is linear, and the underlying prior distribution is Gaussian [33]. For nonlinear forward models, however, the same choice of α_n leads to instabilities that are often fixed with ad-hoc methods such as covariance localization. The main objective of the present work is to show numerically that such instabilities can be addressed by the proposed ensemble Kalman method when the ensemble size N_e is sufficiently large.

In this subsection we briefly present numerical evidence of the regularization properties of the proposed scheme; a detailed numerical investigation will be presented in Section 4. We apply Algorithm 1 to the identification of the “true” log hydraulic conductivity u^\dagger displayed in Figure 1 (top). We use synthetic measurements of hydraulic head from the Darcy model of subsection 2.3.1 with the measurement locations displayed in Figure 3 (left). Both the truth and the elements of the initial ensemble are generated from a Gaussian distribution. For details on the generation of synthetic data (avoiding inverse crimes) and the generation of initial ensemble, we refer the reader to subsection 4.1. In the middle row (resp. bottom row) of Figure 1 we display the estimate obtained from the ensemble mean at different iterations computed with the standard unregularized approach (resp. the proposed regularized method). For the regularized method we select $\rho = 0.7$ and $\tau = 1/\rho$. For the unregularized method we simply apply (19) with no stopping criterion. Both methods are applied with the same (fixed) initial ensemble of $N_e = 150$ members generated as described in subsection 4.1. Figure 2 (left) shows the data misfit (3) for the ensemble mean \bar{u}_n computed with the standard unregularized approach (solid red line) and our regularized method (dotted-black line). Additionally, in Figure 2 (left) we also display (the dotted-blue line) the data misfit with

respect to the averaged ensemble data predictions (expression (6)) which we monitor for the termination of Algorithm 1.

In Figure 2 (middle) we show the relative L^2 -error with respect to the truth for the ensemble mean \bar{u}_n , (i.e. $\|\bar{u}_n - u^\dagger\|_{L^2}/\|u^\dagger\|_{L^2}$), computed with the standard unregularized approach (solid red line) and our regularized method (dotted-blue line). For the standard method we see that a fast decrease of the data misfit and relative error are observed at the early iterations. While the data misfit keeps decreasing, the error starts to increase when apparently the data misfit drops below the noise level whose value is indicated with the horizontal line in Figure 2 (left). This increase in the error comes as no surprise since the corresponding estimates overfit the data. Iterative methods applied for the solution of ill-posed inverse problems often display such behavior [12, 23]. In fact, a similar semiconvergent behavior was reported in the ensemble Kalman method of [16] and inspired the present work where not only a regularization of the estimates need to be introduced (here by means of the parameter α_n) but also the proper early termination of the scheme that avoids fitting the noise. Indeed, we note that the regularizing algorithm (see Figure 2), both relative error and data misfit decreases very slowly. Once the (averaged) data misfit (of expression (6)) is close to the noise level, the algorithm is stopped producing a considerably more accurate estimate of the truth than the one obtained with the unregularized algorithm which blows up the update at the early iterations. However, if the proposed regularized algorithm is not stopped, the error with respect to the truth could potentially increase as we observe in subsequent experiments (see Figure 9). From the log-conductivity estimates (Figure 1 middle and bottom rows), we see that the unregularized method produces uncontrolled estimates of the unknown while the proposed method results in small incremental transitions which eventually captures the truth more accurately. As we will discuss in Section 4, the ability of the proposed method for regularizing the computations of the inverse problem depends on the size of the ensemble N_e .

The selection of α_n according to (34) is crucial for the regularizing of the proposed method. We monitor numerically these values of α_n obtained from the application of Algorithm 1 to the identification problem described in the preceding paragraph. We apply the method for several choices of the parameter ρ ($\rho = 0.4, 0.5, 0.6, 0.7, 0.8$). In Figure 2 (right) we display the values of α_n that we obtained from the proposed selection (34) as the number of iterations increases. At the early iterations α_n is large hence controlling the updates of the approximation. As the iteration progresses and the data misfit decreases, α_n also decreases. For larger ρ the decay of α_n is slower; this provides more regularization to the expense of a more costly algorithm (see Remark 3.1). From Figure 2 (right) we can observe that when the data misfit has dropped to a value close to the noise level by which the algorithm will be terminated according to (39), α_n has dropped down to a value close to $\alpha_n = 1$ which is, in turn, the intrinsic value for the standard Kalman method. Interestingly, there is indeed an apparent intrinsic parameter $\alpha_n = 1$ once the scheme arrives at the optimal approximation of the inverse problem. However, choosing this parameter at early iterations can be substantially

detrimental to the performance of the scheme. In the following section, the selection of α_n is derived by using ensemble approximations of the Fréchet derivative of the forward map in the selection of the parameter α_n that has been proven to provide stability in the regularizing LM scheme of Hanke [14]. As stated earlier, for the present methodology, such a theory is beyond the scope of the present work.




Figure 1. Top: true log-conductivity. Middle row: ensemble mean at some iterations of the standard (unregularized) ensemble Kalman method. Bottom row: ensemble mean at some iterations of the regularizing ensemble Kalman method.

3. Properties and computational aspects of the regularizing ensemble Kalman method

In this subsection we discuss some properties and general computational aspects of the proposed regularizing ensemble Kalman method presented in subsection 2.2.

3.1. The initial ensemble

One of the main properties of the proposed ensemble Kalman method is that the estimate \bar{u}_n obtained from Algorithm 1 lives in the subspace generated by the initial ensemble $\{u_0^{(j)}\}_{j=1}^{N_e}$.

Proposition 3.1 (Invariance Subspace Property) *At every iteration of the scheme, the ensemble mean \bar{u}_n defined by (4) satisfies $\bar{u}_n \in \text{span}\{u_0^{(j)}\}_{j=1}^{N_e}$*




Figure 2. Log-data misfit (left) and relative error w.r.t truth (middle) of the ensemble mean at some iterations obtained with the unregularized ensemble Kalman method (solid red line) and with the proposed regularizing scheme (dotted blue line). Right: $\log(\alpha_n)$ as a function of the number of iterations of Algorithm 1 from 40 experiments with different initial ensembles of size $N_e = 150$.

Proof: Once we rewrite Algorithm 1 in the augmented version of subsection 3.4, the proof follows directly with the same argument as [16, Theorem 2.1] \square .

From the invariance subspace property it follows that the selection of the initial ensemble is a design parameter crucial to the performance of the proposed scheme; we study this numerically in subsection 4.2. Prior knowledge of the space of admissible solutions X can be used for such selection. For example, $\{u_0^{(j)}\}_{j=1}^{N_e}$ can be a truncated basis for X . Another example of prior knowledge that we may use for the construction of the initial ensemble is the regularity of the elements in X . More concretely, we may construct an ensemble by drawing its members from some probability distribution with the desired regularity. While the problem under consideration is deterministic, the use of a probability distribution is made for the sake of the generation of the initial ensemble with the regularity of the parameter space. The aforementioned invariance subspace property then ensures that the estimate produced by the proposed method inherits the regularity of the space of admissible solutions. Clearly, Gaussian distributions are desirable since sampling from them is relatively easy. Nonetheless, other priors such as Besov [25] could also be considered. In subsections 4.2 we provide examples where probability distributions are considered in order to generate an initial ensemble that we use with our computational approach for the EIT problem described in subsection 2.3.2. In general, the construction of the prior ensemble is based on prior knowledge of the problem under consideration.

3.2. The Regularizing LM scheme

In the following subsection we derive the proposed method as an approximation of the regularizing LM scheme [14]. For the subsequent derivation we consider X completed with the norm $\|C^{-1/2} \cdot\|_X$ where $C^{-1} : D(C^{-1}) \subset X \rightarrow X$ is a densely-defined unbounded self-adjoint operator with compact resolvent and $C^{-1/2}$ is defined in terms of

the spectral decomposition of C^{-1} (see [16, Section 2.1]). For the purpose of this work, C^{-1} is an operator selected a priori that enforces regularity on the functional space X . Introducing an operator C^{-1} in our formulation enable us to (i) derive the proposed method from the regularizing LM scheme of Hanke [14] and (ii) establish a connection between the classical (deterministic) and the Bayesian approach for inverse problems. In the Bayesian framework, C^{-1} is the inverse of the covariance operator C from a prior distribution [32]. It is important to remark that the definition of C^{-1} does not appear in the proposed scheme Algorithm 1.

The regularizing LM scheme of [14] is an iterative method that possess the regularizing properties need for the stable computation of solution to the identification problem described in the preceding section. In concrete Hanke in [14] proposes an iterative scheme where the u_{n+1} iteration level is given by

$$u_{n+1} = u_n + \arg \min_{v \in X} \|\Gamma^{-1/2}(y^\eta - G(u_n) - DG(u_n)v)\|_Y^2 + \alpha_n \|C^{-1/2}v\|_X^2 \quad (20)$$

where $\alpha > 0$ is a regularization parameter chosen below and DG denotes the Fréchet derivative of G . From the first order optimality conditions associated to the right hand side of (20) it follows that the previous expression is equivalent to

$$u_{n+1} = u_n + (DG^*(u_n) \Gamma^{-1} DG(u_n) + \alpha_n C^{-1})^{-1} C DG^*(u_n)(y^\eta - G(u_n)). \quad (21)$$

The convergence and regularizing properties of the method of [14] requires that α_n in (20) satisfies

$$\rho \|\Gamma^{-1/2}(y^\eta - G(u_n))\|_Y \leq \alpha_n \|\Gamma^{-1/2}(y^\eta - G(u_n) - DG(u_n)(u_{n+1} - u_n))\|_Y \quad (22)$$

for $\rho \in (0, 1)$ selected a priori and that the scheme is terminated whenever the n th iteration level satisfies

$$\|\Gamma^{-1/2}(y^\eta - G(u_n))\|_Y \leq \tau \eta < \|\Gamma^{-1/2}(y^\eta - G(u_{n-1}))\|_Y \quad (23)$$

for some $\tau > 0$ that satisfies $\tau > 1/\rho$.

The theory of Hanke provides assumptions on the forward operator under which the regularizing LM scheme terminates after a finite number of iterations and the corresponding approximation is a solution of the inverse problem in the small noise limit as described in subsection 2.1. For full details of the theoretical framework for the regularizing LM scheme the reader is referred to the work of Hanke in [14]. An application of this method to inverse problems in the geosciences can be found in [20].

In order to derive the proposed ensemble method as an approximation of the LM scheme, we need the following lemma where we assume that $Y = \mathbb{R}^M$ (recall that $\dim(Y) < \infty$) with the standard Euclidean inner product.

Lemma 3.1 *Assume that for any $u \in X$, the linear functionals $D_m G(u) : X \rightarrow \mathbb{R}$ ($m = 1, \dots, M$) are linearly independent. Then, expression (21) is equivalent to*

$$u_{n+1} = u_n + C DG^*(u_n)(DG(u_n) C DG^*(u_n) + \alpha_n \Gamma)^{-1}(y - G(u_n)) \quad (24)$$

Proof: See appendix.

Note that (24) can be used to rewrite (22) as follows

$$\rho \|\Gamma^{-1/2}(y^\eta - G(u_n))\|_Y \leq \alpha_n \|\Gamma^{1/2}(DG(u_n) C DG^*(u_n) + \alpha_n \Gamma)^{-1}(y^\eta - G(u_n))\|_Y \quad (25)$$

Expressions (24)-(25) are computationally more convenient since the operator inversion of $(DG(u_n) C DG^*(u_n) + \alpha_n \Gamma)$ is conducted on a finite dimensional space. Upon discretization the aforementioned inversion has often negligible cost since the number of observations is typically small. Note that when X is finite dimensional the relation between (21) and (24) follows simply from matrix Lemmas which cannot be applied in the present case.

3.3. The proposed method as an approximation of the regularizing LM scheme

Let us consider a first order approximation of $G(u_n^{(j)})$ around the ensemble mean, i.e.

$$G(u_n^{(j)}) \approx G(\bar{u}_n) + DG(\bar{u}_n)(u_n^{(j)} - \bar{u}_n)$$

Then,

$$\bar{w}_n \equiv \frac{1}{N_e} \sum_{j=1}^{N_e} G(u_n^{(j)}) \approx G(\bar{u}_n) \quad \text{and} \quad \langle DG(\bar{u}_n)(u_n^{(j)} - \bar{u}_n), w \rangle_Y \approx \langle G(u_n^{(j)}) - G(\bar{u}_n), w \rangle_Y. \quad (26)$$

Define the following covariance operator

$$C_n^{uu}(\cdot) = \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (u_n^{(j)} - \bar{u}_n) \langle u_n^{(j)} - \bar{u}_n, \cdot \rangle_X \quad (27)$$

From (26)-(27) and (7) -(8) we find that

$$\begin{aligned} C_n^{uu} DG^*(\bar{u}_n)v &= \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (u_n^{(j)} - \bar{u}_n) \langle (u_n^{(j)} - \bar{u}_n), DG^*(\bar{u}_n)v \rangle_X \\ &\approx \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (u_n^{(j)} - \bar{u}_n) \langle G(u_n^{(j)}) - G(\bar{u}_n), v \rangle = C_n^{uw}v \end{aligned} \quad (28)$$

and from similar arguments we obtain

$$DG(\bar{u}_n)C_n^{uu}DG(\bar{u}_n)^*v \approx C_n^{ww}v \quad (29)$$

In expression (24) we now replace the following terms and the corresponding approximations

$$u_n \implies \bar{u}_n, \quad (30)$$

$$C DG^*(u_n) \implies C_n^{uu} DG^*(\bar{u}_n) \approx C_n^{uw} \quad (31)$$

$$DG(u_n) C DG^*(u_n) \implies DG(\bar{u}_n)C_n^{uu}DG^*(\bar{u}_n) \approx C_n^{ww} \quad (32)$$

thereby obtaining

$$\bar{u}_{n+1} = \bar{u}_n + C_n^{uw}(C_n^{ww} + \alpha_n \Gamma)^{-1}(y^\eta - \bar{w}_n) \quad (33)$$

that we use as the update formula for the ensemble mean of our iterative scheme. Similarly, from (26) - (32), the selection of α in (25) and the stopping criteria (23) become

$$\rho \|\Gamma^{-1/2}(y^\eta - \bar{w}_n)\|_Y \leq \alpha_n \|\Gamma^{1/2}(C_n^{ww} + \alpha_n \Gamma)^{-1}(y^\eta - \bar{w}_n)\|_Y. \quad (34)$$

and (6), respectively. The sequence defined in (10) satisfies (34). In addition, note that the sample mean obtained from the ensemble updated according to formula (9) satisfies indeed (33). We then obtain the proposed ensemble Kalman scheme presented in Algorithm 1.

3.4. The proposed method as a sequence of linear inverse problems

As state earlier, the proof of convergence of Algorithm 1 is beyond the scope of the present manuscript. Nonetheless, in this subsection we show some relevant properties which shed light on its regularizing effect. In order to study the aforementioned properties it is advantageous to rewrite the proposed algorithm in terms of the following augmented variables [16],

$$z = \begin{pmatrix} u \\ w \end{pmatrix}, \quad \Xi(z) = \begin{pmatrix} u \\ G(u) \end{pmatrix}, \quad (35)$$

the space $Z \equiv X \times Y$ and the projection operators H and H^\perp defined by $H z = w$ and $H^\perp z = u$, respectively. It is not difficult to see that, in terms of these new variables, Algorithm 1 becomes

Algorithm 2

Let $\{u_0^{(j)}\}_{j=1}^{N_e} \subset X$ be the initial ensemble of N_e elements. Let $\rho \in (0, 1)$ and $\tau > 1/\rho$. Define

$$z_0^{(j,a)} \equiv \begin{pmatrix} u_0^{(j)} \\ G(u_0^{(j)}) \end{pmatrix} \quad (36)$$

For $n = 1, \dots$

(1) **Prediction step.** Evaluate the forward map (35),

$$z_n^{(j,f)} = \Xi(z_{n-1}^{(j,a)}). \quad (37)$$

and define sample mean:

$$\bar{z}_n^f = \frac{1}{N_e} \sum_{j=1}^{N_e} z_n^{(j,f)} \quad (38)$$

(2) **Discrepancy principle.** If

$$\|\Gamma^{-1/2}(y - H \bar{z}_n^f)\|_Y \leq \tau \eta \quad (39)$$

stop. Output

$$\bar{u}_n \equiv \frac{1}{N_e} \sum_{j=1}^{N_e} H^\perp z_n^{(j,f)} = \frac{1}{N_e} \sum_{j=1}^{N_e} u_n^{(j)} \quad (40)$$

Define sample covariance:

$$C_n z = \frac{1}{N_e - 1} \sum_{j=1}^{N_e} (z_n^{(j,f)} - \bar{z}_n^f) \langle z_n^{(j,f)} - \bar{z}_n^f, z \rangle_Z. \quad (41)$$

(3) **Analysis step.** Update each ensemble member as follows

$$z_n^{(j,a)} = z_n^{(j,f)} + C_n H^* (H C_n H^* + \alpha \Gamma)^{-1} (y^\eta - H z_n^{(j,f)}) \quad (42)$$

where α satisfies

$$\rho \|\Gamma^{-1/2} (y^\eta - H \bar{z}_n^f)\|_Y \leq \|\Gamma^{-1/2} (y^\eta - H \bar{z}_n^a(\alpha))\|_Y \quad (43)$$

with

$$\bar{z}_n^a = \frac{1}{N_e} \sum_{j=1}^{N_e} z_n^{(j,a)} \quad (44)$$

It is not difficult to see [16] that C_n in (41) can be written as

$$C_n = \begin{pmatrix} C_n^{uu} & C_n^{uw} \\ C_n^{wu} & C_n^{ww} \end{pmatrix}$$

While the regularization properties of the proposed method will be studied numerically in Section 4, it is clear that the essence of the proposed regularizing Kalman method is the selection of α by means of (43). Since $0 < \rho < 1$, the existence of such α is ensured by the following proposition

Proposition 3.2 *At every iteration of the scheme, the ensemble mean $\bar{z}_n^a(\alpha)$ defined by (44) satisfies*

- (i) $\|\Gamma^{-1/2} (y - H \bar{z}_n^a(\alpha))\|_Y \rightarrow \|\Gamma^{-1/2} (y - H \bar{z}_n^f)\|_Y$ as $\alpha \rightarrow \infty$.
- (ii) $\|\Gamma^{-1/2} (y - H \bar{z}_n^a(\alpha))\|_Y \rightarrow 0$ as $\alpha \rightarrow 0$.
- (iii) The map $\alpha \rightarrow \|\Gamma^{-1/2} (y - H \bar{z}_n^a(\alpha))\|_Y$ is monotonously nondecreasing

Proof: The proof can be carried out as in [23, Theorem 2.16] \square .

Similar to the regularizing LM scheme where each iterate solves a linear inverse problem (see expression (20)), our ensemble Kalman scheme can be posed as Tikhonov-type regularization. This equivalence motivates the need to further regularize standard Kalman methods.

Note that the ensemble mean of the analysis step (42) is given by

$$\bar{z}_n^a = \bar{z}_n^f + C_n H^* (H C_n H^* + \alpha \Gamma)^{-1} (y^\eta - H \bar{z}_n^f) \quad (45)$$

We now show that (45) can be posed as the solution of a regularized linear inverse problem

Proposition 3.3 *Assume that at each iteration, the ensemble $\{z_n^{(j,f)}\}_{j=1}^{N_e}$ is linearly independent. Then, the ensemble mean of the analysis step (45) of the ensemble Kalman algorithm satisfies*

$$\bar{z}_n^a = \operatorname{argmin}_{z \in \mathcal{Z}_n} \left(\|\Gamma^{-\frac{1}{2}} (y^\eta - H \bar{z}_n^f - H(z - \bar{z}_n^f))\|_Y^2 + \alpha \|z - \bar{z}_n^f\|_{C_n}^2 \right) \quad (46)$$

where \mathcal{Z}_n is the completion of $\mathcal{R}(C_n)$ with respect to the norm induced by $\langle \cdot, C_n^{-1} \cdot \rangle_Z$ and denoted by $\|\cdot\|_{C_n}$.

Proof: From the definition of C_n in (41) it is clear that the rank of C_n is given by $\mathcal{R}(C_n) = \text{span}\{z_n^{(j,f)}\}$. Also note that C_n is a sum of rank-one operators, hence compact. If $N_e < \infty$, then $\dim(\mathcal{R}(C_n)) = N_e < \infty$; thus from the linearly independence of the initial ensemble it follows that the restriction $C_n : \mathcal{R}(C_n) \rightarrow \mathcal{R}(C_n)$ is a positive definite operator. Hence its inverse $C_n^{-1} : \mathcal{R}(C_n) \rightarrow \mathcal{R}(C_n)$ exists. Moreover, we may consider the restriction of H into $\mathcal{R}(C_n)$ as well as its adjoint $H^* : \mathcal{R}(C_n) \rightarrow Y$. Thus, the operators C_n , C_n^{-1} and H can be represented with matrices and so standard matrix algebra can be directly applied to show that (20) is equivalent to [33]

$$(H^*\Gamma^{-1}H + \alpha C_n^{-1})(\bar{z}_n^a - \bar{z}_n^f) = H^*\Gamma^{-1}(y - H\bar{z}_n^f) \quad (47)$$

which, from simple standard arguments is the solution to the normal equations associated to the minimization of

$$J(z) \equiv \|\Gamma^{-\frac{1}{2}}(y - H\bar{z}_n^f - H(z - \bar{z}_n^f))\|^2 + \alpha\|z - \bar{z}_n^f\|_{C_n}^2 \quad (48)$$

in the space \mathcal{Z}_n . For the case $N_e \rightarrow \infty$ we may consider $C_n^{-1} : \mathcal{R}(C_n) \rightarrow \mathcal{R}(C_n)$ as a densely defined unbounded operator. Then a proof based on representers can then be carried out similarly to the proof of Lemma 3.1. \square .

From Proposition 3.3 we notice that the ensemble mean at each iteration of the ensemble Kalman algorithm is the Tikhonov-regularized solution (in the subspace \mathcal{Z}_n with norm $\|\cdot\|_{C_n}$) of the following “artificial” linear inverse problem:

$$\text{given } \tilde{y} \equiv y^\eta - H\bar{z}_n^f \quad \text{find } w \equiv z - \bar{z}_n^f \in \mathcal{Z}_n \quad \text{such that } \tilde{y} = Hw \quad (49)$$

Let us define $z^\dagger = (u^\dagger, G(u^\dagger))^T$ and $w^\dagger \equiv z^\dagger - \bar{z}_n^f$ where u^\dagger , we recall, is the truth. If at the n th iteration we were to solve the original identification problem (i.e find/approximate u^\dagger) that would imply to find w^\dagger defined above. In other words, w^\dagger is the solution to the inverse problem (49) and so the corresponding exact data is $\tilde{y}^\dagger \equiv Hw^\dagger$ which is

$$\tilde{y}^\dagger \equiv Hw^\dagger = H(z^\dagger - \bar{z}_n^f) = G(u^\dagger) - H\bar{z}_n^f = y^\eta - \xi - H\bar{z}_n^f = \tilde{y} - \xi \quad (50)$$

where we have used (1) and the definition of H . From the previous expression and (2) it follows that

$$\|\Gamma^{-1/2}(\tilde{y}^\dagger - \tilde{y})\| = \|\Gamma^{-1/2}\xi\|_Y = \eta \quad (51)$$

which implies that the noise level for the linearized inverse problem solved by \bar{z}_n^a , is the same as the noise level of the original nonlinear inverse problem that we aim at solving by means of the proposed method. Moreover, note that before convergence is achieved, i.e. when $\tau\eta < \|\Gamma^{-1/2}(y^\eta - H\bar{z}_n^f)\|_Y$ from (43) we then have that

$$\eta < \frac{1}{\tau} \|\Gamma^{-1/2}(y^\eta - H\bar{z}_n^f)\|_Y \leq \rho \|\Gamma^{-1/2}(y - H\bar{z}_n^f)\|_Y \leq \|\Gamma^{-1/2}(y^\eta - H\bar{z}_n^a(\alpha))\|_Y. \quad (52)$$

Thus, at each iteration, the selection of α in (43) honors the discrepancy principle for the artificial inverse problem that the ensemble mean solves at each iteration of the scheme.

Remark 3.1 *From the augmented version of the proposed ensemble Kalman method we may appreciate more clearly the role of the parameter ρ . Indeed, note from Proposition*

3.2 that a value of $\rho \approx 1$ ($\rho < 1$) in (43) will result in a very large α and so the estimate $\bar{z}_n^a(\alpha)$ will move slightly from the previous estimate $\bar{z}_n^f(\alpha)$. Then, a selection of $\tau \approx 1/\rho$, ($\tau > 1/\rho$) for $\rho \approx 1$ will enable us to allow the algorithm to progress until the data misfit is close to the noise level. In contrast, a smaller ρ results in a smaller α and thus a less controlled estimate $\bar{z}_n^f(\alpha)$. Therefore, choosing $\tau \approx 1/\rho$, ($\tau > 1/\rho$) for smaller ρ implies that we need to stop the algorithm via (39) much earlier to prevent the data overfitting as previously discussed. While it seems clear that a choice of $\rho \approx 1$ will result in a more stabilized scheme, it may be also detrimental to the performance of the scheme. In Section 4 we investigate, with numerical experiments, practical choices for the parameter ρ .

3.5. The connection with the Bayesian framework

While here we propose an ensemble Kalman method as a derivative-free regularizing method for classical inverse problems, most Kalman-based approaches are posed in a Bayesian inference framework where the objective is to produce an ensemble from which statistical information from the Bayesian posterior can be computed. In order to fully understand the underlying motivation of our method it is instructive to consider the role of the proposed method in the context of Bayesian inversion. Let us then assume that the noise ξ in (1) is centered Gaussian noise with covariance Γ . Additionally, assume that there is an underlying Gaussian prior distribution $N(\bar{u}, C)$ of the unknown u , where the mean and covariance of such distribution are \bar{u} and C respectively. Choose an initial ensemble generated from samples of this distribution, i.e. $u_0^{(j)} \sim N(\bar{u}, C)$. Let us now consider the first iteration $n = 0$ of Algorithm 1 with a fixed parameter $\alpha = 1$ and replace y^η in (9) by $y^{(j)} = y^\eta + \eta^{(j)}$ where $\eta^{(j)} \sim N(0, \Gamma)$. Then, Algorithm 1 becomes

$$u_1^{(j)} = u_0^{(j)} + C_0^{rw}(C_0^{rw} + \Gamma)^{-1}(y^{(j)} - G(u_0^{(j)})) \quad (53)$$

which is also the standard formula for the so-called Ensemble Smoother [13]. If the forward operator G is linear then it can be shown [16] that the updated ensemble $\{u_1^{(j)}\}_{j=1}^{N_e}$ fully characterizes, as $N_e \rightarrow \infty$, the conditional probability of u given the data y^η . In other words, (53) provides an ensemble approximation of the Bayesian posterior. Moreover, the mean of the ensemble converges, in the limit $N_e \rightarrow \infty$, to the maximum a posteriori estimate, i.e. the maximizer of the aforementioned posterior distribution. The linear-Gaussian case is trivial in the sense that the posterior is Gaussian and it can be easily characterized with its mean and covariance. However, when the forward operator G is nonlinear, such as the ones described in subsection 2.3, the Bayesian posterior is in general non-Gaussian and the ensemble smoother provides only an approximation whose convergence theory, to our best knowledge, is nonexistent. It has been often reported, however, that when straight forward application of (53) are used in the general nonlinear case, inaccurate approximations of the Bayesian posterior may be obtained [21]. In a recent publication [19], we have demonstrated that the application of iterative regularization methods like the ones proposed here may improve the accuracy of ensemble methods for approximation of the statistical inverse problem. The present

work follows the classical approach; we endow an intrinsically Bayesian method with the regularizing properties needed to address deterministic inverse problem in a derivative-free easy-to-implement computational framework.

3.6. Small ensemble size ill-conditioning

As stated above, issues of stability of ensemble Kalman methods for data assimilation have been often reported. Ad-hoc fixes such as covariance localization and covariance inflation have become standard practice to address such lack of stability which have been typically associated and more noted with implementations of small ensemble size with respect to the number of measurements [1]. Indeed, note that since $\dim(Y) = M < \infty$, then $\dim(\mathcal{R}(C_n^{ww})) = \min\{N_e, M\}$. Clearly, when $N_e < M$, C_n^{ww} is singular and, even though $(C_n^{ww} + \Gamma)^{-1}$ is strictly positive definite, it could potentially have a large conditioning number for some choices of Γ , which is in turn selected according to measurement information. This type ill-conditioning when $N_e < M$ has the similar effect of the ill-posedness of the Fréchet derivative of the forward map in the regularizing LM scheme. Thus, it comes as no surprise that the application of ideas from iterative regularization techniques can be useful for addressing the ill-conditioning of the method when used with a small ensemble size with respect to the number of measurements. From either sources of ill-conditioning, it is clear that ensemble Kalman methods need regularization and, to our best knowledge, the standard fixes indicated above do not possess a mathematical framework which provides general guidelines for developing robust implementations. In contrast, iterative regularization has a rigorous general mathematical ground; our goal here is to use it as the tool to empower ensemble Kalman methods with the stabilization/regularization required to applied them as accurate robust but computationally tractable solvers for large-scale PDE-constrained inverse problems.

3.7. Computational aspects of the scheme

Under our assumption that the number of measurements is small compared to the dimension of the (discretized) parameters space, the main computational cost of the proposed scheme is in the evaluation of the forward map in the prediction step (5). The construction of C_n^{rw} , C_n^{ww} as well as the inversion of $(C_n^{ww} + \alpha_n \Gamma)^{-1}$ are negligible for the number of measurement considered for the present applications. Therefore, the cost of Algorithm 1 is mainly $n^* N_e$ forward model evaluations where n^* is the stopping iteration determined by (43). It is clear that Algorithm 1 is computationally feasible provided that, with a few N_e ensemble members, a reasonably accurate estimate can be achieved in only a few iterations. In Section 4 we provide an extensive numerical investigation of the performance and computational feasibility of the scheme. In some cases, a relatively large number of ensembles are required to achieve stable computations with reasonable accuracy.

4. Numerical Experiments

By means of numerical experiments, in this section we investigate the convergence and regularizing properties of the iterative regularizing ensemble Kalman method. We consider the identification problems on the test models from subsection 2.3 and conduct synthetic experiments where we apply the proposed scheme with synthetic data generated with a true parameter. Algorithm 1 then produces an estimate (the ensemble mean (4)) which we compare against the truth in order to assess the accuracy of the scheme. By conducting synthetic experiments, our aim is to study the accuracy, convergence and regularizing properties of the proposed methods with respect to

- ensemble size N_e
- tunable parameters ρ and τ
- number of measurements M
- noise level η
- selection of initial ensemble $\{u_0^{(j)}\}_{j=1}^{N_e}$

4.1. Test Model I. Darcy flow

We now describe the generation of synthetic data from the Darcy flow model of subsection 2.3.1. We consider the identification of the log of the “true” hydraulic conductivity u^\dagger of Figure 3 (left). The true conductivity is a random field drawn from a Gaussian measure $N(0, C)$ with covariance $C : L^2(D) \rightarrow L^2(D)$ defined by

$$C\phi(x) = \int_D c(x, y)\phi(y) dy \quad (54)$$

with a spherical covariance function

$$c(x, y) = \begin{cases} c_0 \left[1 - \frac{3}{2} \frac{h}{a} + \frac{1}{2} \frac{h^3}{a^3} \right] & \text{if } h \equiv |x - y| < a \\ 0 & \text{if } h > a \end{cases} \quad (55)$$

where $a > 0$ is the correlation length and $c_0 > 0$ is the variance of the field. Gaussian distributions with covariance operators of this type are often used to model the geologic properties of some formations [9]. The true log conductivity field is used in equations (11)-(13) to find the true head field displayed in Figure 3 (middle). The latter is then used in (14) where the measurement locations $\{x_m\}_{m=1}^M$ will be specified below. Then, Gaussian noise is added to the synthetic data; the norm of the noise is 1% of the norm of the data. The Darcy flow model is solved numerically with cell-centered finite differences [29]. The matrix Γ is simply a diagonal matrix with entries proportional to the size of the corresponding entry of $G(u^\dagger)$. In order to avoid inverse crimes, Algorithm 1 is applied on a coarser grid (of 80×80 cells) than the one (of 160×160 cells) used for the generation of the synthetic data.

For the experiments of this section, we apply Algorithm 1 with an initial ensemble that consists of N_e (to be defined below) random fields drawn from the same Gaussian measure $N(0, C)$ that we use to generate the truth. Some members of the initial

ensemble are displayed in Figure 4. The generation of these fields (as well as the truth) is conducted by means of the Karhunen-Loeve (KL) expansion of a random field distributed according to $N(0, C)$. Although the truth and the prior ensemble are samples from the same probability distribution, note that (i) the truth is generated on a much finer grid and (ii) the truth (or its projection on the coarser grid) is not part of the initial ensemble.

The motivation for selecting the truth as a realization from the same distribution from the initial ensemble is twofold. On the one hand, in subsurface applications there is often prior knowledge of the geologic properties in terms of a prior probability distribution for the truth. It is then natural, as in the Bayesian approach, to use this distribution to generate the initial ensemble. On the other hand, by choosing the truth and the initial ensemble from the same probability distribution we expect to reduce (although not completely eliminate) the effect that the selection of an initial ensemble has on the performance of the proposed scheme. More precisely, from the invariance subspace property (Proposition 3.1) we know that the approximation provided by our method lives in the subspace generated by the initial ensemble. Therefore, each initial ensemble will produce a different estimate. However, by generating the initial ensemble from the same distribution that we use to generate the truth, our initial ensembles, while different from the truth, have the same regularity and spatial structure (e.g. correlation) of the true field. In the following subsection we study an inverse problem where the truth is a prescribed field with no association whatsoever to the initial ensemble and there we study the effect that the spatial correlation of the initial ensemble has on the accuracy of the proposed ensemble method to identify the truth.

While the initial ensemble generated as described above will provide estimates with the same spatial features of the truth, from the aforementioned subspace property we certainly expect that the results from Algorithm 1 will vary with the selection of the ensemble. Therefore, in order to fully assess the numerical properties of this algorithm without the influence of the initial ensemble, for the present work we will conduct multiple experiments corresponding to different choices of the initial ensemble and provide averages over these experiments and thus reliable quantities related to the average performance of the scheme. It is worth reiterating that the initial ensemble is a design parameter that can be chosen in various ways according to available information.

4.1.1. Effect of the ensemble size N_e . We generate synthetic data as described below with the array of 100 measurement locations displayed in Figure 3 (right). We use these data in Algorithm 1 with a (fixed) parameter $\rho = 0.7$. In order to appreciate the effect of the early termination of the scheme, the algorithm is allowed to progress even when the data misfit goes below η/ρ (recall the stopping criteria (6) requires $\tau > 1/\rho$). In the top (resp. bottom) of Figure 5 we plot the relative error with respect to the truth $\|u_n - u^\dagger\|_{L^2(D)}/\|u^\dagger\|_{L^2(D)}$ (resp. log - data misfit) from 40 different experiments corresponding to different selection of the initial ensemble (recall each initial ensemble is a set of N_e draws from a the Gaussian distribution defined above). Different panels




Figure 3. Left: true log-conductivity. Middle: log-data misfit. Right: relative error with respect to the truth




Figure 4. Some elements from the initial ensemble.

corresponds to different selections of ensemble size N_e . Each of the blue curves (resp. red curves) represents the log-data misfit (resp. error with respect to the truth) that we obtain from the ensemble mean \bar{u}_n (expression (4)) computed with Algorithm 1 initialized with each of the 40 initial ensembles mentioned earlier. We reiterate that by “data misfit” we mean the misfit between data and the average of the model outputs from the ensemble defined in (6) which is, in turn, the quantity monitored for the convergence of the scheme.

The green curve in the bottom (resp. top) of Figure 5 represents the log-data misfit (resp. error w.r.t truth) averaged, at each iteration, over the 40 experiments from different initial ensembles. The dotted vertical line defines the iteration number after which the averaged relative error w.r.t. the truth starts increasing. The dotted horizontal line in Figure 5 (bottom) indicates the log of the value of η/ρ . The solid red line indicates the value of the noise level η . As we expected, these results confirm that Algorithm 1 reduces the relative error w.r.t. truth. However, after some number of iterations, this error will start increasing unless the algorithm is stopped. The results from Figure 5 reveals that there is a critical ensemble size for which, on average (over several experiments with different initial ensembles), the discrepancy principle (6), with $\tau \approx 1/\rho$ is a reliable stopping criteria which terminates the algorithm before the error w.r.t. the truth increases due to data overfitting. In other words, the stability that the

scheme inherits from the regularizing LM scheme depends on the ensemble size. For this experiment the critical ensemble size is $N_e = 125$. In average, for $N_e < 125$ the error with respect to the truth of the ensemble mean will increase before the data misfit has dropped below η/ρ . For $N_e \geq 125$ we observe that the error will increase when the data misfit has reached the value in η/ρ as we would expect if the regularization properties were inherited from the regularizing LM scheme. For larger ensemble sizes $N_e \geq 200$ the relative error with respect to the truth will not increase even when the data misfit takes values below η/ρ but a slow increase starts showing when the data misfit drops below the noise level.

We also note from Figure 5 that, as we increase the ensemble size, the ensemble mean \bar{u}_n provides a more accurate approximation of the truth (provided the scheme is properly stopped). However, a further increase in the ensemble size results in higher computational cost of the scheme. In Figure 6 we show the ensemble mean that we obtain, when Algorithm 1 is initialized with different ensemble sizes N_e and the scheme stopped according to (6) with $\tau = 1/\rho$. For this result, the members of the smaller initial ensembles are contained in the larger ones. We can clearly appreciate that for smaller ensemble sizes, the stabilization of the scheme may be lost and thus its accuracy.




Figure 5. Relative error w.r.t. the truth (top) and log - data misfit (bottom) from 40 different experiments with $\rho = 0.7$ associated to different initial ensembles of size (from left to right) $N_e = 75, 100, 125, 200, 400$

4.1.2. Effect of the number of measurements. We now investigate the relation between the number of measurements and the critical size that enable us to observe the regularization properties inherited by the regularizing LM scheme. As discussed in subsection 2.4 stability issues for $N_e < M$ may arise unless the parameter α_n controls the inversion in (9). We recall that the selection of α_n according to (34) is inspired by the regularizing LM scheme which addresses a different type of ill-posedness (although is reflected in a similar fashion). Thus, this selection of α_n may not necessarily address the stability issues that arise in the ensemble method when $N_e < M$. In the following




Figure 6. Log - conductivity estimates obtained from an initial ensemble size of (from left to right) 75, 100, 150, 200 (the larger ensembles contains the smaller ones). These estimates are the ensemble mean obtained from Algorithm 1 with $\tau = 1/\rho$ in (6).

experiments we show that under some conditions, the selection of α_n and the stopping criteria does indeed resolve the stability issues even when $N_e < M$.

We conduct a set of experiments, each of them similar to the one described in the preceding subsection but with synthetic data generated from different measurement configurations. For each measurement configuration and ensemble size, we conduct a set of 40 experiments corresponding to different initial ensembles generated as described earlier. Then, log-data misfit and relative error with respect to the truth were computed and averaged (over the 40 experiments) at each iteration of the scheme. For clarity we only report these averaged quantities corresponding to each ensemble size and each measurement configurations. In the left column of Figure 7 and Figure 8 we display the measurement configuration for each experiment. The middle and right columns of Figures 7 and 8 show the (averaged over different experiments) log data misfit and error w.r.t truth for four different ensemble sizes (increasing from top to bottom in the legend of these Figures). The dotted horizontal line in the middle columns of Figure 7 and Figure 8 denotes the value of the data misfit corresponding to η/ρ . The vertical line in the middle and right columns indicate the iteration after which the error w.r.t truth increases for the second choice (from top to bottom) of ensemble size displayed on the legend of these figures. It is important to note that, irrespective of the ensemble size, the proposed algorithm does stabilize the estimates obtained in the early iterations. In other words, it avoids uncontrolled estimates that some standard unregularized Kalman-based methods may display in the first couple of iterations. However, the relation between ensemble size and number of measurements has severe effects on the ability of the stopping criteria (6) to successfully terminate the algorithm under the guidelines that arise from the application of iterative regularization (i.e. with $\tau \approx 1/\rho$). Let us, for

example, examine the case $M \leq 100$ reported in Figure 7. For smaller measurement locations (i.e. $M = 25, 36$), an ensemble size equal to the number of measurements $N_e = M$, on average, will provide an estimate whose error with respect to the truth will decrease in the first iterations but then starts increasing long before the data-misfit has reached the value η/ρ . A similar but less drastic behavior is observed for larger number of measurements ($M = 49, 64$). However, as we increase M , we find that a selection $N_e = M$ yields results that seemed to be stabilized whenever the scheme is stopped according to (6) with $\tau = 1/\rho$. Note that for these measurement configurations with $M \leq 100$, the ensemble size needs to be larger than the number of measurements so that the computational solution is properly stabilized with the proposed stopping criteria.

Let us consider now the case $M > 100$ reported in Figure 8. In contrast to the previous case where an ensemble of a size larger than that number of measurements is needed for the stabilization of the computations, here certain choices of N_e with $N_e < M$ may suffice. Note for example that for $M = 225$, an ensemble with $N_e = 169$ will result in stable computations by using (6) with $\tau \approx 1/\rho$. Interestingly, for larger M it is more evident that the selection $N_e > M$ is unnecessary for the algorithm to exhibit the regularization properties based on the stopping criteria mentioned above. Indeed, for $M = 400$ and $N = 484$ we note that $N_e = 196$ and $N_e = 225$ will clearly provide regularized estimates in the sense that their error will not increase before the data misfit reaches the value η/ρ . For these cases with $M > 100$, a selection of $N_e \geq M$ will clearly result in not only stable but an accurate identification. However, it is worth reiterating that for large-scale applications, the smallest ensemble size that produce a regularized solution is desirable due to the high computational cost of the forward simulations. In summary, there seems to be an apparent critical number of measurements ($M = 100$) below of which a selection of $N_e > M$ is needed for the proposed method to fully stabilize the scheme. However, for larger number of measurements the regularizing properties of the scheme enable us to use a selection of $N_e < M$ which results in a reasonable computational cost.

4.1.3. The tunable parameters ρ . As we discussed in subsection 3.4, the parameter ρ controls the ensemble updates. The choice of this parameter has a significant effect on both accuracy and cost of the proposed scheme. We now investigate this effect with our Darcy flow model with the measurement configuration of 100 locations displayed in Figure 3 (right). The results from the preceding section indicate that an ensemble of size $N_e > 100$ will provide the stabilization needed with the choice $\tau \approx 1/\rho$ in the stopping criteria. We select $N_e = 150$ for conducting a set of experiments where Algorithm 1 is applied with several choices of ρ . For each choices of ρ , the experiment is repeated 40 times with a different selection of the initial ensemble generated as we previously described. In Figure 9 we display, for these 40 experiments, the log-data misfit (bottom) and the relative error with respect to the truth (top) of the ensemble mean obtained with the proposed scheme with different choices of ρ . The horizontal dotted line in Figure 9 (bottom) represent the log of η/ρ . Note that the stopping




Figure 7. Numerical results with the proposed method for $\rho = 0.7$ and different measurement configurations and ensemble sizes. Left column: measurement configuration. Middle column: relative error w.r.t. the truth. Right column: Log - data misfit. Quantities are averaged at each iteration, over 40 experiments corresponding to different experiment.




Figure 8. Numerical results with the proposed method for $\rho = 0.7$ and different measurement configurations and ensemble sizes. Left column: measurement configuration. Middle column: relative error w.r.t. the truth. Right column: Log - data misfit. Quantities are averaged at each iteration, over 40 experiments corresponding to different experiment.

criteria (6) with $\tau \approx 1/\rho$ provides a reasonable stabilization of algorithm. For $\rho < 0.7$ we observe that the error with respect to the truth increases quite rapidly when the data misfit drops below the value η/ρ . We note that, on average, there is a slight decrease of the error with respect to the truth as we increase ρ . However, this slight increase in the accuracy of the identification has associated an increase in the computational cost. For the Darcy flow model, our experiments suggest that $\rho = 0.7$ represents a reasonable choice in terms of accuracy and cost.




Figure 9. Regularizing ensemble method applied with $N_e = 150$, 100 measurements and ρ (from left to right) $\rho = 0.5, 0.6, 0.7, 0.8, 0.85$. Top row: relative error w.r.t. truth. Bottom row: Log - data misfit.

4.1.4. Regularizing properties of the proposed method in the small-noise limit In this section we conduct experiments to study the regularization properties of the proposed scheme in the small noise limit. The previous experiments show that with a reasonable selection of the ensemble size with respect to the measurement configurations, the mean of the proposed ensemble method achieve stable approximations of the truth. In other words, the early termination of the scheme, say at iteration n^* , yields an approximation $u^\eta \equiv \bar{u}_{n^*}$ given by the mean of the ensemble at the stopped iteration. As discussed earlier, the aim of our iterative regularizing algorithm is to provide a stable approximation in the sense that, as $\eta \rightarrow 0$, $u^\eta \rightarrow u$ where $G(u) = G(u^\dagger)$. Due to the lack of uniqueness of the identification problem we cannot claim that $u = u^\dagger$ (the estimate converges to the truth). Nevertheless, we expect that the elements $u \in X$ such that $G(u) = G(u^\dagger)$ will possess similar spacial features of the truth.

We consider again the configuration with $M = 100$ as before and we vary the ensemble size. In Figure 10 we show the relative error with respect to the truth (top) and the log-data misfit (bottom) obtained for each ensemble size with a different selection of noise levels. The level of noise is selected so that the norm of the noise relative to the data is of the percentage indicated in the plots. The matrix Γ as selected as described earlier but kept fixed for these experiments associated to different noise levels. For this particular experiment the algorithm is stopped according to (6) with $\tau = 1/\rho$. We note

that for the smallest noise considered here 0.5% an ensemble of size $N_e = 150$ was needed to properly stabilize the computations. Again, this shows that the proposed method, for sufficient large ensemble, inherits the regularization properties of the regularization LM scheme of [14].




Figure 10. Log - data misfit (bottom) and relative error w.r.t. the truth (top) from 40 experiments with different initial ensembles of size of (from left to right) 75, 100, 150, 200 and different choices of noise level.

4.1.5. Comparison with the regularizing LM scheme. In Figure 11 we display the numerical results from the application of the proposed scheme to identify the log-conductivity described in the preceding subsections. As before, we used 40 different experiments corresponding to different initial ensembles of size $N_e = 150$. In this case, however, we compare with the regularizing LM scheme applied to the solution of the parameter identification but constrained to the subspace generated by the initial ensemble. The numerical results displayed in Figure 11 suggests that the proposed ensemble Kalman method produces a derivative-free approximation of the regularizing LM scheme where the Fréchet derivative is approximated by the ensemble covariance. Moreover, these results show that the proposed ensemble Kalman algorithm is minimizing the data misfit in a stable fashion.

4.2. Test Model II: EIT

In this section we investigate the performance of the regularizing ensemble Kalman method for the solution of the EIT problem described in subsection 2.3.2. Similar to the previous subsection, our aim is to understand fundamental aspects of the regularizing properties of the scheme. In particular, in this subsection our aim is to observe the effect of the selection of the initial ensemble on the accuracy and cost of the proposed




Figure 11. Log-data misfit (left) and rel. error w.r.t truth (right) from estimates obtained with the regularizing LM scheme and the proposed ensemble method (Results from 50 experiments with different initial ensembles of size $N_e = 150$).

technique. In contrast to the previous subsection where the true unknown field was a random draw from a distribution that we used as well for the generation of the initial ensemble, here we prescribe a conductivity κ that consist of three circular inclusions on a circular domain D of unit radius similar to the one used in [28]. This true conductivity is displayed in Figure 12 along with the configuration of 16 electrodes used in the CEM described in subsection 2.3.2. We consider 15 current patterns where current is injected between a pair of adjacent electrodes. For each current pattern we collect measurements on all the electrodes thereby having $M = 240$ observations. We choose a value $z_k = 0.01$ for the contact impedances of all the electrodes.

In the middle and middle-right of Figure 12 we show some of the true voltages obtained with the FEM solver from EIDORS MATLAB framework [2]. Synthetic data were generated from the aforementioned voltages by adding Gaussian random noise of standard deviation of 2% of the signal. Inverse crimes are avoided by using a finer mesh (with 7744 elements) than the one used for the application of Algorithm 1 (mesh of 6400 elements). The experiments in this subsection are focused on the electrode and measurement configuration described earlier.

Note that the experiments from this subsection comprise the more general case where there is no characterization of the truth in terms of a probability distribution. In this case, we still consider an initial ensemble generated from samples from a Gaussian measure $N(0, C)$. However, this probability distribution is completely artificial and defined only for the purpose of generating an initial ensemble. In other words, we assume no link between the truth and our choice of initial ensemble. We consider the covariance operator that arises from the Wittle-Matter correlation function [28]. In 2D, such covariance operator takes the form

$$C = \omega L^2 (1 - L^2 \Delta)^{-\theta} \quad (56)$$




Figure 12. Left: true log-conductivity. Middle and Right: Two of the 15 voltage simulations computed with the true conductivity.

where Δ is the Laplacian operator, L is a correlation length. For the purpose of this work we regard θ as a parameter that controls the regularity of the samples and ω as a scaling constant. Some samples from such distribution with parameter $L = 0.2$, $\theta = 5$ and $\omega = 0.1$ are displayed in Figure 13. As before, samples from $N(0, C)$ are obtained by means of the KL expansion of random fields $N(0, C)$. Samples are generated on a regular rectangular domain that contains the computational domain shown in Figure 12 and projected on the FEM domain used for the computation of the CEM.




Figure 13. Draws from a Gaussian prior of the form (56) with $L = 0.2$ and $\theta = 5$.

4.2.1. Ensemble size N_e , tunable parameter ρ and small noise limit. In subsection 4.1 we have extensively studied the role of ensemble size and tunable parameters. The aim here is to briefly verify that similar results are obtained when Algorithm 1 is applied to the EIT problem. For a fixed parameter $\rho = 0.5$ we conduct a set of experiments where the scheme is applied to 40 different initial ensembles generated from the Gaussian measure described above. The corresponding log-data misfit (bottom) and relative error (top) from these experiments are displayed in Figure 14. As with the previous test model, the ensemble size is crucial to the performance of the proposed iterative scheme. We observe a critical ensemble size of $N_e = 75$ above which the proposed scheme is properly stabilized when the early termination is carried out according to (6) with the selection of $\tau \approx 1/\rho$ provided by the theory of the regularizing LM scheme. In this case we note again that the critical size above which the method exhibits regularizing properties was smaller than the number of measurements; this confirms that the proposed scheme addresses the small ensemble effect described in subsection 2.4. The ensemble mean \bar{u}_n

from 5 experiments (from different initial ensembles) with $N_e = 100$ are displayed in Figure 15.

In Figure 16 we fix $N_e = 100$ and now perform a set of experiments with different selections of ρ in (34). As before, we confirm that for this sufficiently large ensemble size, the selection of $\tau > 1/\rho$ in the termination of the scheme ensures stabilized estimates. It is worth noticing that the aforementioned selection of τ is essential as the relative error increases abruptly once the data misfit drops below the value η/ρ . While there is a slight increase in the accuracy (in terms of the error w.r.t truth) as we increase ρ , we clearly observe that the number of iterations and so the computational cost of the scheme increases substantially. For this case we observe that $\rho = 0.5$ provides a reasonable balance between computational cost and accuracy.

Finally, in Figure 17 we show the effect of the stabilization of the scheme as the noise decreases. These results corresponds to the application of the scheme with synthetic data generated from the electrode configuration described before but with different noise levels. The results from Figure 17 corresponds to averages, at each iteration, from 40 experiments obtained from 40 different selections of initial ensembles. As we expected, the proposed ensemble scheme inherits the regularizing properties of the LM method and produces stable computations that converge, in a stable fashion as $\eta \rightarrow 0$.




Figure 14. Relative error w.r.t. truth (top) and log - data misfit (bottom) obtained from Algorithm 1 with (from left to right) $N_e = 50, 60, 75, 150, 250$. These plots display the results from 40 experiments with different initial ensembles.

4.2.2. The selection of the initial ensemble As we have indicated earlier, the proposed scheme is highly dependent on the selection of the initial ensemble. In the experiments from the previous subsection, we attempt to reduce such dependence by selecting the initial ensemble from the same distribution that we use to generate the truth. This selection on the initial ensemble, from Proposition 3.1 ensures that the estimate from Algorithm 1 has the same regularity/spatial structure of the truth. We wish now to investigate the effect of the selection of the initial ensemble on the estimates produced




Figure 15. Top: True log-conductivity. Middle and bottom: Estimates of log-conductivity obtained from 5 experiments with different initial ensembles. The proposed method is used with $\rho = 0.5$ and $N_e = 100$.




Figure 16. Error w.r.t. truth (top) and log - data misfit (bottom) obtained with Algorithm 1 with (from left to right) $\rho = 0.3, 0.5, 0.7, 0.8$ from 40 experiments with different initial ensembles of size $N_e = 100$

by our method. We therefore require a systematic way to generate substantially different initial ensembles. This can be achieved, for example, by modifying the parameters L and θ in (56) that, in turn, control the spatial correlation and regularity of the samples that we use as members of such initial ensemble. For simplicity we focus only on the selection of L and consider initial ensembles with the covariance (56) for different choices of L . In Figure 18 we present some samples generated with several choices of L (from left to right: $L = 1, 0.2, 0.1, 0.06, 0.05$). Each row corresponds to a fixed set of realizations of coefficients in the KL expansion that we use to generate those Gaussian fields. We can observe visually how the spatial correlation of the samples decreases with L .

We apply Algorithm 1 for the solution of the EIT problem with an ensemble of size $N_e = 100$ generated from the Gaussian measure described earlier with different




Figure 17. Log data misfit (right) and error w.r.t truth (left) from the average over 40 experiments where Algorithm 1 was applied with different initial ensembles to synthetic data of noise levels of 1%, 5%, 7.5%, 10% 15%.




Figure 18. Draws from a Gaussian prior of the form (56) with (from left to right) $L = 1, 0.2, 0.1, 0.06, 0.05$. Each row corresponds to a different realization of the KL coefficients in the parameterization of the Gaussian field.

choices of parameter L . For each value of L , we report the results from 40 experiments corresponding to different initial ensembles (i.e. different draws from the Gaussian measure). Figure 19 displays the resulting log-data misfit (bottom) and error w.r.t truth (top). The selection of the initial ensemble based on L has a significant effect on the performance of the method. We note that the accuracy degrades as we decrease the correlation length. For the largest correlation length considered here ($L = 1$) we note that after a few iterations, before the data misfit reaches the value η/ρ , the relative error shows a slight increase which results in large fluctuations of the data misfit. As we decrease the correlation length the relative error is stabilized and the data misfit

reduced. However, we note that the accuracy decreases as we use initial ensembles generated with members that have smaller correlation lengths ($L < 0.1$). Indeed, for $L = 0.05$, the relative error with respect to the truth increases before the data misfit drops below η/ρ . It is thus clear that the selection of the initial ensemble (in this case based on L) affects the regularizing properties of the scheme.

In Figure 20 we display estimates (i.e. ensemble mean \bar{u}_n) of the log-conductivity obtained from one of the the previous experiments for the different choices of L . More concretely, these are the estimates obtained with different initial ensemble corresponding to different choices of L but generated with the same realization of KL coefficients. From the invariance subspace property of the proposed scheme, it comes as no surprise that an initial ensemble of elements that have very small spatial correlation cannot possibly produce an estimate that characterizes the truth. Indeed, the inclusions of high conductivity from the truth have a diameter of approximately 0.3 which can be better identified when the initial ensemble is generated from smooth fields with a similar spatial correlation. Similar dependence on correlation lengths were obtained in the Bayesian level-set approach of [17]. While prior knowledge of the truth can be certainly used for the construction of the initial ensemble, further research should address the potential identification of parameters that determine the regularity of the initial ensemble that better characterize the truth.




Figure 19. Error with respect to the truth (top) and log data misfit (bottom) from 40 experiments with different initial ensembles generated from samples of $N(0, C)$ with values of L in (56) (from left to right) $L = 0.2, 0.1, 0.06, 0.04, 0.03$.

5. Applications to geometric inverse problems

The numerical investigations of the preceding sections establishes that the proposed ensemble Kalman method, with a sufficiently large ensemble size, inherits the regularizing properties of the regularizing LM approach of [14]. Thus, our scheme can be




Figure 20. Top: Estimates of level-set obtained from 5 experiments with different initial ensembles with samples of $N(0, C)$ with values of L in (56) (from left to right) $L = 0.2, 0.1, 0.06, 0.05$. Bottom: True conductivity.

potentially applied to produce stable and accurate estimates of PDE-constrained inverse problems. Moreover, those estimates can be computed at a reasonable computational cost without the need of the derivative of the forward map. In this section, we show the potential application of the proposed method to solve identification problems where the computation of such derivative can be very cumbersome. In concrete, in this section we investigate the application of the proposed regularizing ensemble Kalman method for the solution of shape identification problems.

We are interested in the identification of an unknown region Ω within D , the spatial domain of definition of the PDE under consideration, whose boundary $\partial\Omega$ defines a sharp discontinuity of an unknown parameter in the PDE model. That is the case, for example, in subsurface flow applications where the conductivity of the aquifer/reservoir may take substantially different (by several orders of magnitude) values on facies characterized with different geologies. While the nominal values on such geologic facies may be known a priori, the interface between geologic facies is usually unknown.

A common approach to parameterize an unknown interface (or shape) is to use an implicitly parameterization in terms of level-sets [26]. In other words, we may assume that

$$\Omega = \{x \in D \mid u(x) \leq 0\} \quad (57)$$

where here u denotes the level-set function. The zero level-set provides the unknown interface which is now parameterized by a function u and thus the identification can be posed within a functional-analytical framework. Level-set-based methods for shape identification problems have been extensively applied in the literature [6, 5, 18, 30]. However, to our best of knowledge, most of these approaches use a variational framework for solving shape-based least-squares problems which, in turn, requires computation of the shape derivatives of the forward map. In addition, these standard variational formulations require the solution of the level-set equation to evolve the shape so that it minimizes the, possibly regularized, squared data misfit. This approach often leads

to computations of flat level-set functions that need to be redefined. In this section we show that the proposed ensemble Kalman method can be used to approximate the solution to identification problems without using neither the shape derivative of forward maps nor the level-set equation. More precisely, we apply the proposed iterative ensemble Kalman method for the solution of identification problems where the interface is parameterized by (57). The update formulate (9) with the controlled/regularized selection of α induces the motion of the shape. Moreover, we propose to apply the Kalman method with a selection of the initial ensemble of level-sets that consist of samples from Gaussian measures $N(0, C)$ which a covariance that enforces some smoothness of the level-set function and could potentially incorporate some intrinsic correlation length. Nonetheless, this probabilistic approach for the generation of the initial ensemble is conducted artificially for the sake of the implementation. However, it has been recently shown [17] that there is zero probability of generating samples (and thus initial ensembles) corresponding to flat level-set functions. This, combined with the invariance subspace property of the proposed method, ensures that our estimates do not become flat thus overcoming the common issue of the flattening of the level-set function typical of standard level-set approaches.

5.1. Estimation of geologic facies

We are interested in the identification of geologic facies in the test model described in section 2.3. More precisely we want to estimate the region of high conductivity Ω (represented by (57)) in an aquifer D . For this example we prescribe a true conductivity K^\dagger displayed in Figure 21 (middle). This true conductivity simulates a typical layered structure in the geologic properties of an aquifer. We use K^\dagger in (11)-(13) to generate synthetic data that we wish to invert in order to identify the true conductivity. The measurement configuration is displayed in Figure 21 (right). As before, inverse crimes are avoided by using a finer grid for the generation of the truth than the one used for the initial ensemble and so for the inversion.

Our goal is now to apply the proposed ensemble Kalman method on the variable u corresponding to the level-set function that parametrizes κ as follows

$$\kappa(u) = \kappa_i \chi_{u \leq 0} + \kappa_e \chi_{u > 0}, \quad (58)$$

where χ_A denotes the characteristic function of region A and κ_i and κ_e are the positive constants that define the high and low conductivity values of the true conductivity Figure 21 (left). We assume these two constants are known; the unknown is the interface between the regions of different conductivity.

As described earlier we prescribe an artificial Gaussian measure from which we generate members of an initial ensemble of level-set functions. More concretely we consider samples from $N(0, C)$ with

$$C = \omega(-\Delta)^{-\theta} \quad (59)$$

where Δ is the Laplacian operator, ω is a scaling constant and θ controls the regularity. In contrast to (56), the previous expressions defines a covariance with an intrinsic

fixed correlation length that we are not able to vary. In Figure 22 (top row) we show some samples from the prior distribution and in Figure 22 (bottom row) we show the corresponding conductivities obtained from (58).




Figure 21. Left: True hydraulic conductivity. Middle: true hydraulic head. Right: Measurement locations.




Figure 22. Top: Different realizations from level-set functions sampled of $N(0, C)$ with C from (59). Bottom: Conductivities computed from these sample level-set functions (by means of (58)).

We apply the regularizing ensemble Kalman method proposed with different ensemble sizes ($N_e = 75, 100, 150, 250, 350$) and with a fixed parameter $\rho = 0.7$. As before, we consider 40 experiments with different initial ensembles and the corresponding log-data misfit and error with respect to the truth are displayed in Figure 23. Although we are applying the ensemble algorithm to the estimation of the level-set function u , we consider the relative error with respect to the truth of the corresponding conductivity given by (58). Note that a large ensemble is needed in order to obtain stable computations in the scheme until the data misfit reaches the value η/ρ that we use to stop the algorithm. This value is indicated with the dotted horizontal line in the bottom of Figure 23. For sufficiently large ensemble size ($N_e > 150$), on average both data misfit

and error decrease. Stable computations are obtained when the algorithm is stopped via (6) with $\tau \approx 1/\rho$. In the top and middle rows of Figure 24 we display estimates of the level-set function and the corresponding conductivity obtained with Algorithm 1 (for $N_e = 150$) from five experiments with different initial ensemble sizes (with initial ensemble generated from the same Gaussian measure). The visual agreement between the estimates of conductivity and truth (Figure 24 bottom) is remarkable.




Figure 23. Error with respect to the truth (top) and log data misfit (bottom) from 40 experiments with different initial ensembles of size (from left to right) $N_e = 75, 100, 150, 250, 350$.

5.2. EIT. Sharp interfaces.

In this subsection we consider again the EIT problem but assume that we are now interested in recovering a conductivity with sharp discontinuities. Let us then consider the conductivity and the electrode configuration displayed in Figure 25 (left). We use this field as the true conductivity that we aim at estimating with the ensemble Kalman method by using the level-set parametrization of expression (58). Synthetic data are generated (and avoiding inverse crimes) as described in subsection 4.2. Some true voltages are displayed in Figure 25 (middle and right panels). Level-set approaches for shape identification in EIT problems have been studied, for example, in [8, 27]. However, these approaches are variational and use a level-set equation for the evolution of the shape.

As in the preceding subsection, we postulate an artificial Gaussian measure $N(0, C)$ that we use to generate an initial ensemble. We consider again C described by (56). However, note that for this subsection, such Gaussian is used to generate the ensemble of level-set function rather than conductivities. Nonetheless, as in the experiments of subsection 4.2, we expect the parameter L to have an influence on the initial ensemble and this the estimate of the level-set and the associated conductivity obtained with (58). In Figure 26 we present some samples of the level-set function




Figure 24. Top: Estimates of level-set obtained from 5 experiments with different initial ensembles with samples of $N(0, C)$ with C from (54). Middle: Conductivities computed from these estimated level-set functions (by means of (58)). Bottom: True hydraulic conductivity.




Figure 25. Top-left: true conductivity. Top-middle and top-right: two of the 15 voltage patterns generated with the true conductivity.

(top) and the corresponding conductivity (bottom) obtained with different values of L ($L = 0.2, 0.1, 0.06, 0.04, 0.03$) but with the same set of KL coefficients. We clearly observe that the correlation length of the level-set is reflected in the spatial correlation of the interface between the regions of different conductivities. In Figure 27 we show estimates of level-set function (top) and conductivities (bottom) obtained, by means of Algorithm 1, with an initial ensemble of $N_e = 200$ elements generated from the Gaussian distribution described above with the aforementioned values of L . We clearly observe that there is an optimal choice of L which yields estimates that visually agree better with the truth displayed in Figure 27 (bottom). In this case, $0.06 \leq L \leq 0.1$ provides

(on average) the lowest error w.r.t. truth. In Figure 28 we show the relative error (top) and the log-data misfit (bottom) from different experiments corresponding to multiple initial ensembles generated with the Gaussian $N(0, C)$ for the aforementioned choices of L . For larger correlation lengths $L \geq 0.2$ the data misfit displays larger fluctuations that seem to arise from the appearance of high conductivity close to the electrodes. In addition, very small values of L yields conductivities where the interface between different values has short correlation length. In the current framework, knowledge of the optimal correlation length can be used for the generation of initial ensemble. However, for more general/realistic cases where such parameter is unknown, the estimation of L should be conducted within the method; this is beyond the scope of the present manuscript.




Figure 26. Top: Level-set functions sampled of $N(0, C)$ with values of L in (56) (from left to right) $L = 0.2, 0.1, 0.06, 0.04, 0.03$ (these samples have the same realization of KL coefficients). Bottom: Conductivities computed from these sample level-set functions (by means of (58)).

6. Conclusions

Our numerical investigation indicates that for sufficiently large ensemble size, the proposed regularizing ensemble Kalman method inherits the regularizing properties of the LM scheme [14]. In other words, when N_e is large enough, we observe that the proposed selection of the regularization parameter and early termination of the scheme prevent the lack of stability typical of unregularized schemes. Fortunately, the aforementioned size needed for stability is reasonable and often used in standard ensemble implementations for large-scale applications. For the Darcy problem, for example, we found that with $N_e = 150$ stable and reasonably accurate estimates can be obtained, in average, within 12 iterations of the scheme when $\rho = 0.7$ (and with $M = 100$ measurements). The computational cost for computing these estimates is 1800 forward model simulations which is comparable with other gradient-based techniques where also dozens of iterations are need for the convergence, but at each iteration, the explicit computation of the Fréchet derivative (if possible at all) has often the cost of M linearized forward model simulations (recall M is the total number of




Figure 27. Top: Estimates of level-set obtained from 5 experiments with different initial ensembles with samples of $N(0, C)$ with values of L in (56) (from left to right) $L = 0.2, 0.1, 0.06, 0.04, 0.03$. Middle: Conductivities computed from these estimated level-set functions (by means of (58)). Bottom: True conductivity.




Figure 28. Error with respect to the truth (top) and log data misfit (bottom) from 40 experiments with different initial ensembles generated from samples of $N(0, C)$ with values of L in (56) (from left to right) $L = 0.2, 0.1, 0.06, 0.04, 0.03$.

measurements). From the results of the preceding section we can clearly appreciate that the computational cost and accuracy of the proposed method with large enough N_e is very competitive with variational iterative regularization methods. However, as we have consistently reiterated throughout this manuscript, the main advantage of the proposed method is that no derivatives of the forward map are needed for the proposed ensemble regularizing scheme.

For the EIT problem, we found that a reasonable ensemble size $N_e = 100$ is sufficient for the method to display stability in the case of 16 electrodes considered

in the subsection 2.3.2 and for wide class of tunable parameters. However, we observe that the selection of the initial ensemble is crucial for the stability and accuracy of the proposed scheme. In concrete, there is a range for the optimal selection of the spatial correlation length of the members of the initial ensemble. As we expected, good approximations of the truth are obtained if the samples that we use for the generation of the initial ensemble have a correlation length similar to the one of the true conductivity. While prior knowledge of the regularity and spatial features from the truth could be incorporated for the selection of the initial ensemble, we recognize that in more general cases such information may not be available. Further research should address the estimation of the parameters that control the regularity/spatial features of the initial ensemble.

The numerical experiments from the preceding section show that the proposed method is computationally flexible and suitable for a wide class of parameterizations of PDE parameters. We demonstrated that the regularizing ensemble Kalman method can be successfully used for the identification of shapes where the geometry is parameterized with a level-set function. The advantage of having a regularized ensemble method has been showcased with the examples of Section 5 where we observe that, with a reasonable ensemble size, the proposed method provides stable computations of the unknown regions of high conductivity. The regularization properties and our selection of the initial ensemble generated from a Gaussian distribution is key for developing a robust level-set scheme that avoids solving the level-set equation for the evolution of the shape; our scheme induces the motion of the shape in a controlled fashion.

Standard ensemble Kalman methods have been numerously applied for solving identification problems where the forward map arises from a large-scale forward model. However, these standard implementations are often unstable in particular when small ensemble sizes (like the ones used here) are used. We have shown that the proposed regularizing ensemble method, when the ensemble size is large enough, not only addresses the ill-posedness proper of the inverse problem but also the one associated to the small size effect. There is a broad class of ensemble methods where the Fréchet derivative of the forward map is approximated with an ensemble [10]. While most of these methods have been developed to address concrete applications in a statistical setting, they could be potentially applied for the solution of generic parameter identification problems if these methods are empowered with the regularization/stabilization needed to compute the solution of such inverse problems. This manuscript offers numerical evidence that importing ideas from iterative regularization methods can potentially use with ensemble methods to stabilize/regularize the computations of inverse estimates.

Acknowledgments The author would like to thank Andrew Stuart for his helpful comments on the manuscript. The author also wishes to thank Matthew Dunlop for his help with using EIDORS.

7. Appendix

7.1. Proof of Lemma 3.1

Note that if X is a finite dimensional, $DG(u_n)$, $D^*G(u_n)$ and C in (21) are matrices. Moreover, since in this case C is invertible (recall C is positive definite) simple matrix algebra can be applied to show the desired equivalence [33]. When X is an infinite-dimensional space the argument for matrices can no longer be applied. In this case, however, we can use a representer-based argument. First we note that, since $Y = \mathbb{R}^M$, $DG(u_n) : X \rightarrow Y$ has the form $DG(u_n)v = [DG_1(u_n)v, \dots, D_M G(u_n)v]$ with $D_j G(u_n) : X \rightarrow \mathbb{R}$ for all $j \in \{1, \dots, M\}$. Moreover, we note that

$$w DG_m(u_n)v = \langle v, D_m G^*(u_n)w \rangle_X \quad (60)$$

for all $w \in \mathbb{R}$ and $v \in X$. Therefore,

$$\langle w, DG(u_n)v \rangle_{\mathbb{R}^M} = \sum_{m=1}^M w_m D_m G(u_n)v = \left\langle v, \sum_{m=1}^M D_m G^*(u_n)w_m \right\rangle_X$$

and so $DG^*(u_n)w \equiv \sum_{m=1}^M D_m G^*(u_n)w_m$ for all $w \in \mathbb{R}^M$. Let us denote $\{e_i\}_{i=1}^M$ the canonical basis in \mathbb{R}^M . Note that

$$\begin{aligned} DG(u_n) C DG^*(u_n)e_i &= DG(u_n) C DG_i^*(u_n)1 \\ &= [D_1 G(u_n) C DG_i^*(u_n)1, \dots, D_M G(u_n) C DG_i^*(u_n)1] \end{aligned} \quad (61)$$

and so

$$e_j^T DG(u_n) C DG^*(u_n)e_i = D_j G(u_n) C D_i G^*(u_n) \quad (62)$$

Let $r_m \in X$ be the unique “representer” such that

$$D_m G(u_n)v = \langle C^{-1/2}r_m, C^{-1/2}v \rangle_X \quad (63)$$

for all $v \in X$. On the other hand, from (60) we have $D_m G(u_n)v = \langle v, DG_m^*(u_n)1 \rangle_X$ for all $m \in \{1, \dots, M\}$, then $C^{-1}r_m = DG_m^*(u_n)1$ and so, formally,

$$r_m = C DG_m^*(u_n)1 \quad (64)$$

We may now consider the following representation in terms of representer

$$v = \sum_{m=1}^M \beta_m r_m + b \quad (65)$$

where $b \perp r_j$ in X for all $j \in \{1, \dots, M\}$ (i.e. $\langle C^{-1/2}r_m, C^{-1/2}b \rangle_X = 0$). From (64) and (65) we have

$$DG(u_n)v = \sum_{m=1}^M \beta_m DG(u_n)r_m + DG(u_n)b = \sum_{j=1}^M \beta_m DG(u_n)C D_m^* G(u_n)1 + DG(u_n)b \quad (66)$$

Note that $D_m G(u_n)b = \langle C^{-1/2}r_m, C^{-1/2}b \rangle_X = 0$. Therefore, we arrive at

$$DG(u_n)v = \sum_{j=1}^M \beta_m DG(u_n)C D_m^* G(u_n)1 = DG(u_n)C DG^*(u_n)\beta \quad (67)$$

where $\beta = (\beta_1, \dots, \beta_M)$. Similarly, from (62)- (65) we obtain

$$\begin{aligned} \|C^{-1/2}v\|_X^2 &= \sum_{jm} \beta_j \beta_m \langle C^{-1/2}r_m, C^{-1/2}r_j \rangle_X + \langle C^{-1/2}b, C^{-1/2}b \rangle_X \\ &= \sum_{jm} \beta_j \beta_m D_m G(u_n) r_j + \|C^{-1/2}b\|_X^2 = \beta^T DG(u_n) CDG^*(u_n) \beta + \|C^{-1/2}b\|_X^2 \end{aligned} \quad (68)$$

Combining (67) and (68) we write

$$J(v) \equiv \|\Gamma^{-1/2}(y^\eta - G(u_n)) - DG(u_n)v\|_Y^2 + \alpha \|C^{-1/2}v\|_X^2$$

as

$$\begin{aligned} J_{LM}(\beta, b) &\equiv \|\Gamma^{-\frac{1}{2}}(y - G(u_n) - DG(u_n)CDG^*(u_n)\beta)\|^2 \\ &\quad + \alpha \beta^T DG(u_n)CDG^*(u_n)\beta + \alpha \|C^{-1/2}b\|_X^2 \end{aligned} \quad (69)$$

From the linearly independence of $\{D_m G(u_n)\}_{m=1}^M$, it is not difficult to see that the matrix $DG(u_n)CDG^*(u_n)$ is invertible. Then, from standard arguments it follows that the unique minimizer of (69) in $\mathbb{R}^M \times (\text{span}\{r_m\}_{m=1}^M)^\perp$ is

$$\beta = [DG(u_n)CDG^*(u_n) + \alpha \Gamma]^{-1}(y - G(u_n)), \quad b = 0 \quad (70)$$

Expression (24) then follows from writing $u_{n+1} - u_n = v$ and using (64)-(65) and (70). \square

References

- [1] S. I. Aanonsen, G. Nvdal, D. S. Oliver, A. C. Reynolds, and B. Valles. The ensemble kalman filter in reservoir engineering—a review. *SPE Journal*, 14.
- [2] Andy Adler and William R B Lionheart. Uses and abuses of eiders: an extensible software base for eit. *Physiological Measurement*, 27(5):S25, 2006.
- [3] J. Bear. *Dynamics of Fluids in Porous Media*. Dover Publications, New York, 1972.
- [4] Liliana Borcea. Electrical impedance tomography. *Inverse Problems*, 18(6):R99, 2002.
- [5] M. Burger. A framework for the construction of level set methods for shape optimization and reconstruction. *Interfaces and Free Boundaries*, 5:301–329, 2002.
- [6] M. Burger and S. Osher. A survey on level set methods for inverse problems and optimal design. *Eur. J. Appl. Math.*, 16:263–301, 2005.
- [7] J. Carrera and S. P. Neuman. Estimation of aquifer parameters under transient and steady state conditions: 3. application to synthetic and field data. *Water Resources Research*, 22.
- [8] Eric T. Chung, Tony F. Chan, and Xue-Cheng Tai. Electrical impedance tomography using level set representation and total variational regularization. *Journal of Computational Physics*, 205(1):357 – 372, 2005.
- [9] C. V. Deutsch. *Geostatistical Reservoir Modeling*. Oxford University Press, Oxford, 2002.
- [10] AlexandreA. Emerick and AlbertC. Reynolds. Investigation of the sampling performance of ensemble-based methods with a simple reservoir model. *Computational Geosciences*, 17(2):325–350, 2013.
- [11] H W Engl, K Kunisch, and A Neubauer. Convergence rates for Tikhonov regularisation of non-linear ill-posed problems. *Inverse Problems*, 5(4):523, 1989.
- [12] H.W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*, volume 375. Springer, 1996.
- [13] G. Evensen. *Data assimilation: the ensemble Kalman filter*. Springer Verlag, 2009.
- [14] M. Hanke. A regularizing Levenberg-Marquardt scheme, with applications to inverse groundwater filtration problems. *Inverse Problems*, 13:79–95, 1997.

- [15] M. A. Iglesias and C. Dawson. The representer method for state and parameter estimation in single-phase Darcy flow. *Computer Methods in Applied Mechanics and Engineering*, 196(1):4577–4596, 2007.
- [16] Marco A Iglesias, Kody J H Law, and Andrew M Stuart. Ensemble kalman methods for inverse problems. *Inverse Problems*, 29(4):045001, 2013.
- [17] Marco A Iglesias, Yulong Lu, and Andrew Stuart. A bayesian level-set method for geometric inverse problems. *submitted*, 2015.
- [18] Marco A Iglesias and Dennis McLaughlin. Level-set techniques for facies identification in reservoir modeling. *Inverse Problems*, 27(3):035008, 2011.
- [19] MarcoA. Iglesias. Iterative regularization for ensemble data assimilation in reservoir models. *Computational Geosciences*, pages 1–36, 2014.
- [20] MarcoA. Iglesias and Clint Dawson. The regularizing levenbergmarquardt scheme for history matching of petroleum reservoirs. *Computational Geosciences*, pages 1–21, 2013.
- [21] MarcoA. Iglesias, KodyJ.H. Law, and AndrewM. Stuart. Evaluation of gaussian approximations for data assimilation in reservoir models. *Computational Geosciences*, pages 1–35, 2013.
- [22] B. Katltenbacher, A. Neubauer, and O. Scherzer. *Iterative Regularization Methods for Nonlinear Ill-Posed Problems*. Radon Series on Computational and Applied Mathematics, de Gruyter, Berlin, 1st edition, 2008.
- [23] Andreas Kirsch. *An Introduction to the Mathematical Theory of Inverse Problems*. Springer-Verlag New York, Inc., New York, NY, USA, 1996.
- [24] G. Li and A. Reynolds. Iterative Ensemble Kalman Filters for Data Assimilation. *in Proceedings of the SPE Annual Technical Conference and Exhibition, Anaheim, California, USA, 11-14 Noviembre*, SPE 109808, 2007.
- [25] M.Dashti, S. Harris, and A.M.Stuart. Besov priors for bayesian inverse problems. *Inverse Problems and Imaging*, 6:183–200, 2012.
- [26] S. Osher and J.A. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.
- [27] Peyman Rahmati and Andy Adler. A level set based regularization framework for eit image reconstruction. *Journal of Physics: Conference Series*, 434(1):012083, 2013.
- [28] Lassi Roininen, Janne M. J. Huttunen, and Sari Lasanen. Whittle-mate rn priors for bayesian statistical inversion with applications in electrical impedance tomography. *Inverse Problems and Imaging*, 8(2), 2014.
- [29] T.F. Russell and M.F. Wheeler. Finite element and finite difference methods for continuous flows in porous media. *In: R.E. Ewing, Editor, Mathematics of Reservoir Simulation*, SIAM, Philadelphia, PA.
- [30] F. Santosa. A level-set approach for inverse problems involving obstacles. *ESAIM Contrle Optim. Calc. Var*, 1:17–33, 1996.
- [31] Erkki Somersalo, Margaret Cheney, and David Isaacson. Existence and uniqueness for electrode models for electric current computed tomography. *SIAM Journal on Applied Mathematics*, 52(4):pp. 1023–1040, 1992.
- [32] A.M. Stuart. Inverse problems: a Bayesian perspective. In *Acta Numerica*, volume 19. 2010.
- [33] A. Tarantola. *Inverse Problems Theory: Methods for Data Fitting and Model Parameter Estimation*. Elsevier,, New York, 1st edition, 1987.